

INEXACT KRYLOV SUBSPACE ALGORITHMS FOR LARGE MATRIX EXPONENTIAL EIGENPROBLEM FROM DIMENSIONALITY REDUCTION

GANG WU*, TING-TING FENG[†], LI-JIA ZHANG[‡], AND MENG YANG[§]

Abstract. Matrix exponential discriminant analysis (EDA) is a generalized discriminant analysis method based on matrix exponential. It can essentially overcome the intrinsic difficulty of small sample size problem that exists in the classical linear discriminant analysis (LDA). However, for data with high dimension, one has to solve a large matrix exponential eigenproblem in this method, and the time complexity is dominated by the computation of exponential of large matrices. In this paper, we propose two inexact Krylov subspace algorithms for solving the large matrix exponential eigenproblem effectively. The contribution of this work is threefold. First, we consider how to compute matrix exponential-vector products efficiently, which is the key step in the Krylov subspace method. Second, we compare the discriminant analysis criterion of EDA and that of LDA from a theoretical point of view. Third, we establish a relationship between the accuracy of the approximate eigenvectors and the distance to nearest neighbour classifier, and show why the matrix exponential eigenproblem can be solved approximately in practice. Numerical experiments on some real-world databases show superiority of our new algorithms over many state-of-the-art algorithms for face recognition.

Key words. Large matrix exponential eigenproblem, Krylov subspace method, Dimensionality reduction, Face recognition, Linear discriminant analysis (LDA), Exponential discriminant analysis (EDA).

AMS subject classifications. 65F15, 65F10.

Face recognition has become one of the most successful applications of image analysis. Real data for face images are usually depicted in high dimensions. In order to handle high dimensional data, their dimensionality needs to be reduced. In essence, dimensionality reduction is the transformation of high-dimensional data into a lower dimensional data space. Currently, one of the most extensively used dimensionality reduction methods is subspace transformation [12, 41, 42, 44].

Linear discriminant analysis (LDA) is one of notable subspace transformation methods for dimensionality reduction [12, 14, 28]. LDA encodes discriminant information by maximizing the between-class scatter, and meanwhile minimizing the within-class scatter in the projected subspace. More precisely, let $\mathcal{X} = [\chi_1, \chi_2, \dots, \chi_n]$ be a set of training samples in a d -dimensional feature space, and assume that the original data in \mathcal{X} is partitioned into k classes as $\mathcal{X} = [\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_k]$. We denote by n_j the number of samples in the j -th class, and thus $\sum_{j=1}^k n_j = n$. Let μ_j be the centroid of the j -th class, and μ be the global centroid of the training data set. If we

*Corresponding author. G. Wu is with Department of Mathematics, China University of Mining and Technology & School of Mathematics and Statistics, Jiangsu Normal University, Xuzhou, 221116, Jiangsu, P.R. China. E-mail: gangwu76@126.com and wugangzy@gmail.com. This author is supported by the National Science Foundation of China under grant 11371176, the Natural Science Foundation of Jiangsu Province under grant BK20131126, the 333 Project of Jiangsu Province, and the Talent Introduction Program of China University of Mining and Technology.

[†]Department of Mathematics, Shanghai Key Laboratory of Pure Mathematics and Mathematical Practice, East China Normal University, Dongchuan RD 500, Shanghai, 200241, P.R. China. E-mail: tofengtingting@163.com.

[‡]School of Mathematics and statistics, Jiangsu Normal University, Xuzhou, 221116, Jiangsu, P.R. China. E-mail: zhanglijia86@163.com.

[§]School of Computer Science and Technology, Soochow University, Suzhou, 215006, Jiangsu, P.R. China. E-mail: eyangmeng@163.com.

denote $\mathbf{e}_j = [1, 1, \dots, 1]^T \in \mathbb{R}^{n_j}$, then the within-class scatter matrix is defined as

$$S_W = \sum_{j=1}^k \sum_{\chi_i \in \mathcal{X}_j} (\chi_i - \mu_j)(\chi_i - \mu_j)^T = H_W H_W^T, \quad (1.1)$$

where $H_W = [\mathcal{X}_1 - \mu_1 \cdot \mathbf{e}_1^T, \dots, \mathcal{X}_k - \mu_k \cdot \mathbf{e}_k^T] \in \mathbb{R}^{d \times n}$. The between-class scatter matrix is defined as

$$S_B = \sum_{j=1}^k n_j (\mu_j - \mu)(\mu_j - \mu)^T = H_B H_B^T, \quad (1.2)$$

where $H_B = [\sqrt{n_1}(\mu_1 - \mu), \sqrt{n_2}(\mu_2 - \mu), \dots, \sqrt{n_k}(\mu_k - \mu)] \in \mathbb{R}^{d \times k}$. The LDA method is realized by maximizing the between-class scatter distance while minimizing the within-class scatter distance, which involves solving a “Trace Ratio” problem [19, 23, 24, 25] in the form of

$$\tau = \max_{\substack{V \in \mathbb{R}^{d \times t} \\ V^T V = I}} \frac{\text{tr}(V^T S_B V)}{\text{tr}(V^T S_W V)}, \quad (1.3)$$

where $\text{tr}(\cdot)$ denotes the trace of a matrix, and $t \ll d$ is the dimension of projection subspace. However, this problem is seldom solved in practice. It is generally considered too difficult to solve and is commonly replaced by a simpler, but not equivalent, problem called “Ratio Trace” problem of the following form [23, 28]

$$\varrho = \max_{\substack{V \in \mathbb{R}^{d \times t} \\ V^T V = I}} \text{tr}((V^T S_W V)^{-1} (V^T S_B V)), \quad (1.4)$$

and the optimal projection matrix V can be calculated from solving the following generalized symmetric eigenproblem

$$S_B \mathbf{x} = \lambda S_W \mathbf{x}. \quad (1.5)$$

In practice, the dimension of real data usually exceeds the number of training samples, which results in the scatter matrix S_W being singular. This is called the *small-sample-size* (SSS) or *undersampled problem* [28]. It is an intrinsic limitation of the classical LDA method, and it is also a common problem in classification applications [28]. In other words, the SSS problem stems from generalized eigenproblems with singular matrices. To tackle the SSS problem, many variants of LDA have been proposed in recent years. To name a few, the regularized LDA method (RLDA) [15], LDA+PCA [5], the null-space LDA method (NLDA) [8], LDA/QR [49], LDA/GSVD [18, 48], the direct LDA method (DLDA) [50], the orthogonal LDA method (OLDA) [26, 47], the neighborhood minmax projections (NMMP) [27], and so on. The above variations on LDA have both advantages and disadvantages [28, 51]. An *et al.* [2] unified these LDA variants in one framework: principal component analysis plus constrained ridge regression.

Recently, a novel method based on matrix exponential, called exponential discriminant analysis (EDA), was proposed in [51]. Instead of the LDA criterion (1.4), EDA considers the following criterion

$$\rho = \max_{\substack{V \in \mathbb{R}^{d \times t} \\ V^T V = I}} \text{tr}((V^T \exp(S_W) V)^{-1} (V^T \exp(S_B) V)), \quad (1.6)$$

where $\exp(\cdot)$ denotes exponential of a matrix or scalar in this paper. The projection matrix V can be obtained from solving the t dominant eigenvectors of the following *generalized symmetric matrix exponential eigenproblem* [51]

$$\exp(S_B)\mathbf{x} = \lambda \exp(S_W)\mathbf{x}. \quad (1.7)$$

The framework of the EDA method for dimensionality reduction has gained wide attention in recent years. For instance, an exponential locality preserving projections (ELPP) [40] was proposed to avoid the SSS problem occurred in locality preserving projections (LPP). Wang *et al.* [41] applied matrix exponential to extend many popular Laplacian embedding algorithms such as locality preserving projections, unsupervised discriminant projections, and marginal fisher analysis. A matrix exponential local discriminant embedding method (ELDE) was investigated in [11] to deal with the SSS problem appeared in local discriminant embedding (LDE). Using the method of exponential discriminant analysis, Ahmed [1] proposed a novel image clustering model that incorporated both local and global information in image database. A 2DEDA method was presented in [46], which is an algorithm based on image matrices (2D data) rather than image vectors (1D data). Thus, 2DEDA can be viewed as a generalization of the EDA method to 2D data.

It has been widely shown that the EDA method has more discriminant power than its original counterpart [1, 11, 40, 41, 46, 51]. In the EDA framework, the matrix exponential can be roughly interpreted as a random walk over the feature similarity matrix, and thus is more robust. As the exponentials of S_W and S_B are symmetric positive definite (SPD), the EDA method naturally deals with the SSS problem. Moreover, the behavior of the decay property of matrix exponential is more significant in emphasizing small distance pairs [41]. However, in all the EDA-based methods, one has to solve the large matrix exponential eigenproblem (1.7) for data with high dimensionality [1, 11, 40, 41, 46, 51], and the time complexity is dominated by the computation of $\exp(S_B)$ and $\exp(S_W)$ [51, pp.191]. This cost will be prohibitively large as the exponential of a matrix is often dense, even if the matrix in question is sparse [17]. Thus, for data with high dimension, the EDA method often suffers from heavy overhead and storage requirement. So it is urgent to seek new technologies to speed up the solution of the large generalized matrix exponential eigenproblem arising in the framework of EDA.

Modern numerical linear algebra exploits the Krylov subspace method in different advanced iterative procedures for large scale eigenvalue problems [4, 31, 37]. Indeed, this type of method ranks among “*The Top 10 Algorithms of the 20th Century*” [9]. In this paper, we devote ourselves to solving (1.7) with the Krylov subspace method. The key involves the evaluation of matrix exponential-vector products, which is a hot topic in large scientific and engineering computations [17, 22]. In conventional approaches, one has to evaluate these products by using some iterative methods [17, 22]. In this work, we derive closed-form formulae for the matrix exponential-vector products, so that these products can be formed very efficiently. The second contribution of this work is to give a theoretical comparison for the discriminant analysis criteria (1.4) and (1.6) of LDA and EDA, respectively, and show the reason why EDA can improve the classification rate of the original LDA method. Finally, we establish a relationship between the accuracy of the approximate eigenvectors and the distance to the nearest neighbour classifier (NN) [10], and shed light on why the matrix exponential eigenproblem can be solved approximately in practical calculations.

This paper is organized as follows. In Section 2, we briefly overview the EDA method. In Section 3, we propose two inexact Krylov subspace algorithms for EDA.

TABLE 1.1
Notations used in this paper

Notations	Descriptions
d	Data dimension
k	Number of classes
n	Number of samples
t	Dimension of projection subspace (or number of desired eigenpairs)
ℓ	Number of samples in the training set
\mathcal{X}	Training samples in a d -dimensional feature space
\mathcal{X}_i	The i -th class of training samples
χ_i	The i -th sample
S_W	The within-class scatter matrix defined in (1.1)
S_B	The between-class scatter matrix defined in (1.2)
$\exp(S_W), \exp(S_B)$	Exponentials of S_W and S_B
$\exp^{1/2}(-S_W)$	Square root of $\exp(-S_W)$
λ, \mathbf{x}	Eigenvalue and eigenvector
$\mathcal{K}_m(A, \mathbf{v}_1)$	Krylov subspace with respect to A and \mathbf{v}_1
$\text{tr}(A)$	Trace of the matrix A
V	Projection matrix composed of the desired eigenvectors
A^T	Transpose of A
I	The identity matrix
O	Zero matrix or vector
$\text{span}\{W\}$	Subspace spanned by the columns of W
$\dim(W)$	Dimension of the subspace W
$\ \cdot\ _2$	2-norm of a vector or matrix

Numerical experiments on some real-world face recognition sets including AR, CMU-PIE, Extended YaleB, FERET, ORL, and Yale are performed in Section 4. Concluding remarks are given in Section 5. Some notations used are listed in Table 1.1.

2. The exponential discriminant analysis method. Given an arbitrary square matrix A , its exponential is defined as [16, 17]

$$\exp(A) = \sum_{j=0}^{\infty} \frac{A^j}{j!} = I + A + \frac{A^2}{2} + \cdots + \frac{A^s}{s!} + \cdots, \quad (2.1)$$

where I is the identity matrix. An important consequence is that *any* matrix exponential is *invertible*. Indeed, we have

$$\exp(-A) \cdot \exp(A) = \exp(A - A) = I,$$

i.e.,

$$\exp^{-1}(A) = \exp(-A). \quad (2.2)$$

Suppose that $A \in \mathbb{R}^{d \times d}$ is symmetric and let $A = X\Lambda X^{-1}$ be the eigen-decomposition of A , where $\Lambda = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_d\}$ is diagonal and X is a $d \times d$ matrix with its columns being the eigenvectors. Then it is easy to see that

$$\exp(A) = X \exp(\Lambda) X^{-1}. \quad (2.3)$$

In other words, A and $\exp(A)$ share the same eigenvectors, and the eigenvalues of $\exp(A)$ are $\exp(\Lambda) = \text{diag}\{\exp(\lambda_1), \exp(\lambda_2), \dots, \exp(\lambda_d)\}$.

The exponential discriminant analysis method (EDA) makes use of the exponential criterion (1.6), and the projection matrix V is obtained from solving the matrix exponential eigenproblem (1.7). This method is equivalent to transforming the original data into a new space by distance diffusion mapping, and the LDA criterion is applied in such a new space [51]. The EDA algorithm is described as follows:

ALGORITHM 1. The exponential discriminant analysis method (EDA) [51]

Input: The data matrix $\mathcal{X} = [\chi_1, \chi_2, \dots, \chi_n] \in \mathbb{R}^{d \times n}$, where χ_j represents the j -th training image.

Output: The projection matrix V .

1. Compute the matrices S_B , S_W , $\exp(S_B)$, and $\exp(S_W)$;
2. Compute the eigenvectors $\{\mathbf{x}_i\}$ and eigenvalues $\{\lambda_i\}$ of $\exp(S_W)^{-1}\exp(S_B)$;
3. Sort the eigenvectors $V = \{\mathbf{x}_i\}$ according to λ_i in decreasing order;
4. Orthogonalize the columns of the projection matrix V .

As a result of diffusion mapping, the margin between different classes is enlarged in the EDA framework, which is helpful in improving classification accuracy [51]. In the small-sample-size case, EDA can extract not only the most discriminant information that is contained in the null space of the within-class scatter matrix, where it is similar to NLDA; but also the discriminant information that is contained in the non-null space of the within-class scatter matrix, and it is equivalent to LDA+PCA [51]. In addition, there are at most $k-1$ nonzero generalized eigenvalues in LDA, where k is the number of classes in the data set, so the dimension of the projected subspace is at most $k-1$ in LDA. As a comparison, there is no (theoretical) dimensionality limitation in EDA, and we can project the input data into a low-dimensional subspace whose dimension is larger than the number of classes.

Example 1.1 We illustrate the superiority of the EDA method over the LDA method via a toy example. Consider the set of training samples $\mathcal{X} = [\chi_1, \chi_2, \chi_3] \in \mathbb{R}^{3 \times 3}$, whose dimension is equal to the number of samples. Let the original data be partitioned into $k = 3$ classes: $\mathcal{X} = [\chi_1, \chi_2, \chi_3] = [\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3]$, and assume that the columns of \mathcal{X} are linear independent. So we have $\mu_i = \chi_i, i = 1, 2, 3$, the global centroid $\mu = (\mu_1 + \mu_2 + \mu_3)/3$, and

$$H_W = [\chi_1 - \mu_1, \chi_2 - \mu_2, \chi_3 - \mu_3] = O,$$

$$H_B = [\chi_1 - \mu, \chi_2 - \mu, \chi_3 - \mu].$$

Therefore, $S_W = H_W H_W^T = O$, $S_B = H_B H_B^T$, and the generalized eigenproblem (1.5) for the LDA method becomes

$$S_B \mathbf{x} = O. \quad (2.4)$$

Since the columns of \mathcal{X} are linear independent, we have $\text{rank}(S_B) = 2$, and there is only one solution vector (up to a scalar) for (2.4).

On the other hand, the matrix-exponential eigenproblem (1.7) for the EDA method reduces to

$$\exp(S_B) \mathbf{x} = \lambda \mathbf{x}, \quad (2.5)$$

which has three linear independent solution vectors. Therefore, the LDA method loses some useful discriminant information, and the EDA method can perform better than the LDA method.

As was stressed in [51, pp.191], however, the time complexity of EDA is dominated by the computation of $\exp(S_B)$ and $\exp(S_W)$, as well as the evaluation of the large matrix exponential eigenproblem. More precisely, one has to explicitly form and store $\exp(S_B)$ and $\exp(-S_W)$, as well as their product $\exp(-S_W)\exp(S_B)$. The complexity is $\mathcal{O}(d^3)$, which is prohibitively large as the dimension of the data is high [17]. Moreover, we have to solve a large matrix exponential eigenproblem with respect to $\exp(S_W)^{-1}\exp(S_B)$. If the matrix exponential eigenproblem is solved by using the implicit QR algorithm [4, 16], another $\mathcal{O}(d^3)$ flops are required [4, 16]. Therefore, it is necessary to seek new strategies to improve the numerical performance of the EDA method.

3. Inexact Krylov subspace algorithms for matrix exponential discriminant analysis. In this section, we propose two inexact Krylov subspace algorithms for solving (1.7). We first consider how to compute the matrix exponential-vector products involved in the Krylov subspace method, and then derive new lower and upper bounds on the EDA criterion (1.6) and the LDA criterion (1.4). Finally, we provide a practical stopping criterion for solving the matrix-exponential eigenproblem approximately.

3.1. Computing the matrix exponential-vector products efficiently. The generalized matrix exponential eigenproblem (1.7) is mathematically equivalent to the following *standard nonsymmetric matrix exponential eigenproblem*

$$\exp(-S_W)\exp(S_B)\mathbf{x} = \lambda\mathbf{x}, \quad (3.1)$$

where we used $\exp^{-1}(S_W) = \exp(-S_W)$. As it is only required to calculate t ($t \ll d$) dominant eigenpairs of the large matrix $\exp(-S_W)\exp(S_B)$, we are interested in solving the large matrix exponential eigenproblem by using the Krylov subspace method.

The Arnoldi method is a widely used Krylov subspace method for finding a few extreme eigenvalues and their associated eigenvectors of a large nonsymmetric matrix [4, 31, 37]. It requires only matrix-vector products to extract eigen-information to compute the desired solutions, and thus is very attractive in practice when the matrix is too large to be solved by, e.g., using the implicit QR algorithm [4, 16]; or the matrix does not exist explicitly but in the form of being capable of generating matrix-vector multiplications.

The principle of the Arnoldi method is described as follows. Given a real nonsymmetric matrix A and an initial real vector \mathbf{v}_1 of unit norm, in exact arithmetics, the m -step Arnoldi process recursively generates an orthonormal basis $\mathcal{V}_m = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$ for the Krylov subspace

$$\mathcal{K}_m(A, \mathbf{v}_1) = \text{span}\{\mathbf{v}_1, A\mathbf{v}_1, \dots, A^{m-1}\mathbf{v}_1\}.$$

At the same time, the projection of A onto $\mathcal{K}_m(A, \mathbf{v}_1)$ is expressed as an upper Hessenberg matrix $H_m = \mathcal{V}_m^T A \mathcal{V}_m$. Afterwards the Rayleigh-Ritz procedure [4, 31, 37] is applied to compute approximate eigenpairs of A : Let $(\tilde{\lambda}, \tilde{\mathbf{y}})$ be some eigenpairs of H_m , then we construct some approximate eigenpairs $(\tilde{\lambda}, \mathcal{V}_m \tilde{\mathbf{y}})$ of A . Here $\tilde{\lambda}$ is called a Ritz value and $\mathcal{V}_m \tilde{\mathbf{y}}$ is called a Ritz vector.

Thus, denote $A = \exp(-S_W)\exp(S_B)$, the key ingredient of the Arnoldi method is to evaluate the matrix exponential-vector products

$$\mathbf{w} = A\mathbf{v} = \exp(-S_W)\exp(S_B)\mathbf{v} \quad (3.2)$$

for different given vectors \mathbf{v} . When the matrices S_W, S_B are large, Krylov subspace methods are widely used for this type of problem [17, 22]. Generally speaking, there are two classes of Krylov subspace methods for evaluating (3.2) *iteratively* [30]. In the first class of methods, the matrix is projected into a much smaller subspace, then the exponential is applied to the reduced matrix, and finally the approximation is projected back to the original large space [32]. In the second class of methods, the exponential function is first approximated by simpler functions such as rational functions, and then the action of the matrix exponential is calculated [43]. In this case, the core of (3.2) reduces to solving some linear systems with multiple shifts [45].

Thanks to the structure of $\exp(-S_W)$ and $\exp(S_B)$, instead of evaluating (3.2) *iteratively*, we compute it directly by using a *closed-form* formula. More precisely, let

$$H_W = Q_W R_W, \quad H_B = Q_B R_B \quad (3.3)$$

be the (skinny) QR decomposition [16] of the matrices H_W and H_B from (1.1) and (1.2), respectively, where $Q_B \in \mathbb{R}^{d \times k}$, $Q_W \in \mathbb{R}^{d \times n}$ are orthonormal, and $R_B \in \mathbb{R}^{k \times k}$, $R_W \in \mathbb{R}^{n \times n}$ are upper triangular. Denote by

$$R_B = U_B \Sigma_B V_B^T, \quad R_W = U_W \Sigma_W V_W^T \quad (3.4)$$

the singular value decomposition (SVD) [16] of R_B and R_W , respectively, and let

$$\tilde{Q}_B = Q_B U_B, \quad D_B = \Sigma_B^2, \quad (3.5)$$

and

$$\tilde{Q}_W = Q_W U_W, \quad D_W = \Sigma_W^2, \quad (3.6)$$

Then we have the following theorem:

THEOREM 3.1. *Let \tilde{Q}_B^\perp and \tilde{Q}_W^\perp be orthonormal bases for the orthogonal complement of $\text{span}\{\tilde{Q}_B\}$ and $\text{span}\{\tilde{Q}_W\}$, respectively, and let f be a function defined on the spectra of S_W and S_B defined in (1.1) and (1.2), respectively. Then*

$$f(S_B) = \tilde{Q}_B f(D_B) \tilde{Q}_B^T + \tilde{Q}_B^\perp f(O_{(d-k) \times (d-k)}) (\tilde{Q}_B^\perp)^T \quad (3.7)$$

and

$$f(S_W) = \tilde{Q}_W f(D_W) \tilde{Q}_W^T + \tilde{Q}_W^\perp f(O_{(d-n) \times (d-n)}) (\tilde{Q}_W^\perp)^T, \quad (3.8)$$

where $O_{(d-k) \times (d-k)}$ and $O_{(d-n) \times (d-n)}$ are $(d-k) \times (d-k)$ and $(d-n) \times (d-n)$ zero matrices, respectively. In particular, for the matrix exponential, we have that

$$\exp(S_B) = \tilde{Q}_B \exp(D_B) \tilde{Q}_B^T + I - \tilde{Q}_B \tilde{Q}_B^T, \quad (3.9)$$

$$\exp(-S_W) = \tilde{Q}_W \exp(-D_W) \tilde{Q}_W^T + I - \tilde{Q}_W \tilde{Q}_W^T, \quad (3.10)$$

and

$$\exp^{1/2}(-S_W) = \tilde{Q}_W \exp\left(-\frac{1}{2}D_W\right) \tilde{Q}_W^T + I - \tilde{Q}_W \tilde{Q}_W^T, \quad (3.11)$$

where $\exp^{1/2}(-S_W)$ represents the square root of $\exp(-S_W)$.

Proof. We only prove (3.7), and the proof of (3.8) is similar. It follows from (1.1), (1.2) and (3.3)–(3.6) that

$$\begin{aligned} S_B &= H_B H_B^T = (Q_B U_B) \Sigma_B (V_B^T V_B) \Sigma_B (U_B^T Q_B^T) = \tilde{Q}_B D_B \tilde{Q}_B^T \\ &= \begin{bmatrix} \tilde{Q}_B & \tilde{Q}_B^\perp \end{bmatrix} \begin{bmatrix} D_B & \\ & O_{(d-k) \times (d-k)} \end{bmatrix} \begin{bmatrix} \tilde{Q}_B^T \\ (\tilde{Q}_B^\perp)^T \end{bmatrix}. \end{aligned}$$

From the properties of matrix function [17, Theorem 1.13], we obtain

$$\begin{aligned} f(S_B) &= \begin{bmatrix} \tilde{Q}_B & \tilde{Q}_B^\perp \end{bmatrix} \begin{bmatrix} f(D_B) & \\ & f(O_{(d-k) \times (d-k)}) \end{bmatrix} \begin{bmatrix} \tilde{Q}_B^T \\ (\tilde{Q}_B^\perp)^T \end{bmatrix} \\ &= \begin{bmatrix} \tilde{Q}_B f(D_B) & \tilde{Q}_B^\perp f(O_{(d-k) \times (d-k)}) \end{bmatrix} \begin{bmatrix} \tilde{Q}_B^T \\ (\tilde{Q}_B^\perp)^T \end{bmatrix} \\ &= \tilde{Q}_B f(D_B) \tilde{Q}_B^T + \tilde{Q}_B^\perp f(O_{(d-k) \times (d-k)}) (\tilde{Q}_B^\perp)^T. \end{aligned}$$

And (3.9)–(3.11) follow from

$$\tilde{Q}_B^\perp (\tilde{Q}_B^\perp)^T = I - \tilde{Q}_B \tilde{Q}_B^T, \quad \tilde{Q}_W^\perp (\tilde{Q}_W^\perp)^T = I - \tilde{Q}_W \tilde{Q}_W^T,$$

and the fact that $\exp(O) = I$ for any square zero matrix. \square

REMARK 3.1. According to (3.9) and (3.10), for a given vector \mathbf{v} , we can compute the matrix exponential-vector product (3.2) within two steps:

- (i) $\mathbf{v} = \tilde{Q}_B \exp(D_B) (\tilde{Q}_B^T \mathbf{v}) + \mathbf{v} - \tilde{Q}_B (\tilde{Q}_B^T \mathbf{v})$;
- (ii) $\mathbf{v} = \tilde{Q}_W \exp(-D_W) (\tilde{Q}_W^T \mathbf{v}) + \mathbf{v} - \tilde{Q}_W (\tilde{Q}_W^T \mathbf{v})$.

Notice that D_B and D_W are $k \times k$ and $n \times n$ diagonal matrices, respectively. As a result, there is no need to explicitly form and store the large matrices S_B , S_W and their exponentials $\exp(S_B)$, $\exp(-S_W)$ in the EDA framework.

Indeed, the generalized matrix exponential eigenproblem (1.7) can also be reformulated as a standard symmetric eigenvalue problem, say, by using $\exp(S_W)$'s Cholesky decomposition [16]. Unfortunately, calculating the Cholesky decomposition of a large matrix can be very expensive. Rather than using the Cholesky decomposition, we consider $\exp(S_W)$'s square root [17] that is advantageous for both computational and theoretical considerations. By (1.7),

$$[\exp^{-1/2}(S_W) \exp(S_B) \exp^{-1/2}(S_W)] [\exp^{1/2}(S_W) \mathbf{x}] = \lambda [\exp^{1/2}(S_W) \mathbf{x}]. \quad (3.12)$$

Denote $\exp^{-1/2}(S_W)$ the inverse of square root of $\exp(S_W)$, and note that $\exp^{-1/2}(S_W) = \exp^{1/2}(-S_W)$, (3.12) can be rewritten as

$$[\exp^{1/2}(-S_W) \exp(S_B) \exp^{1/2}(-S_W)] \mathbf{y} = \lambda \mathbf{y}, \quad (3.13)$$

and

$$\mathbf{x} = \exp^{1/2}(-S_W) \mathbf{y} \quad (3.14)$$

is the desired solution. Notice that the matrix

$$M = \exp^{1/2}(-S_W) \exp(S_B) \exp^{1/2}(-S_W) \quad (3.15)$$

is a symmetric positive definite (SPD) matrix, and our aim is to compute a few dominate eigenpairs of it. The (symmetric) Lanczos method [4, 31, 37] is a widely used

Krylov subspace method for finding a few extreme eigenvalues and their associated eigenvectors of a large symmetric matrix. The Lanczos method can be viewed as a simplification of Arnoldi's method for the particular case where the matrix in question is symmetric. In this case the Hessenberg matrix H_m becomes symmetric tridiagonal, which leads to a three-term recurrence in the Lanczos process. Similar to the Arnoldi process, in the Lanczos process, we need to compute the matrix exponential-vector products

$$\mathbf{w} = M\mathbf{v} = \exp^{1/2}(-S_W)\exp(S_B)\exp^{1/2}(-S_W)\mathbf{v} \quad (3.16)$$

with different vectors \mathbf{v} . In terms of (3.9), (3.10) and (3.11), this problem can be solved within three steps:

$$(iii) \mathbf{v} = \tilde{Q}_W \exp(-1/2 \times D_W) \tilde{Q}_W^T \mathbf{v} + \mathbf{v} - \tilde{Q}_W \tilde{Q}_W^T \mathbf{v};$$

$$(iv) \mathbf{v} = \tilde{Q}_B \exp(D_B) \tilde{Q}_B^T \mathbf{v} + \mathbf{v} - \tilde{Q}_B \tilde{Q}_B^T \mathbf{v};$$

$$(v) \mathbf{v} = \tilde{Q}_W \exp(-1/2 \times D_W) \tilde{Q}_W^T \mathbf{v} + \mathbf{v} - \tilde{Q}_W \tilde{Q}_W^T \mathbf{v}.$$

Compared with the Arnoldi method for the standard nonsymmetric matrix exponential eigenproblem (3.1), the advantage of the Lanczos method for (3.13) is that one can use a three-term recurrence in the Lanczos procedure for a symmetric eigenproblem [4, 31, 37]. The disadvantage is that one has to evaluate three matrix exponential-vector products (iii)–(v) in each step of the Lanczos procedure, rather than two matrix exponential-vector products (i)–(ii) in each step of the Arnoldi procedure. So a natural question is: which one is better, Arnoldi or Lanczos? We believe that the answer is problem-dependent. Indeed, we cannot tell which one is *definitely better* than the other; refer to the numerical experiments given in Section 4. The pseudo-code for computing (3.2) and (3.16) is listed as follows:

ALGORITHM 2. Computing the matrix exponential-vector product (3.2)/(3.16) in the Arnoldi/Lanczos process

• **Preprocessing:**

1. Given the training set \mathcal{X} , form the matrices H_B and H_W ;
2. Computing the skinny QR decompositions: $H_B = Q_B R_B$, $H_W = Q_W R_W$;
3. Computing SVDs: $R_B = U_B \Sigma_B V_B^T$, $R_W = U_W \Sigma_W V_W^T$;
4. Let $Q_B = Q_B U_B$, $Q_W = Q_W U_W$, $\Sigma_B = \exp(\Sigma_B^2)$;

• **Computing:**

5. Given a vector \mathbf{v} ;

if computing (3.2) in the Arnoldi process

- 6A. Let $\Sigma_W = \exp(-\Sigma_W^2)$;
- 7A. $\mathbf{w} = Q_B^T \mathbf{v}$;
- 8A. $\mathbf{v} = Q_B \Sigma_B \mathbf{w} + \mathbf{v} - Q_B \mathbf{w}$;
- 9A. $\mathbf{w} = Q_W^T \mathbf{v}$;
- 10A. $\mathbf{w} = Q_W \Sigma_W \mathbf{w} + \mathbf{v} - Q_W \mathbf{w}$;

end

if computing (3.16) in the Lanczos process

- 6L. Let $\Sigma_W = \exp(-1/2 \times \Sigma_W^2)$;
- 7L. $\mathbf{w} = Q_W^T \mathbf{v}$;
- 8L. $\mathbf{v} = Q_W \Sigma_W \mathbf{w} + \mathbf{v} - Q_W \mathbf{w}$;
- 9L. $\mathbf{w} = Q_B^T \mathbf{v}$;
- 10L. $\mathbf{v} = Q_B \Sigma_B \mathbf{w} + \mathbf{v} - Q_B \mathbf{w}$;
- 11L. $\mathbf{w} = Q_W^T \mathbf{v}$;
- 12L. $\mathbf{w} = Q_W \Sigma_W \mathbf{w} + \mathbf{v} - Q_W \mathbf{w}$;

end

REMARK 3.2. In practical implementations, the “Preprocessing” phase 1–4 is run once for all, and the variables Q_B, Q_W, Σ_B and Σ_W can be stored for a latter use. In this procedure, one needs to perform two skinny QR decompositions in $\mathcal{O}(dk^2)$ and $\mathcal{O}(dn^2)$ flops [16], respectively. Further, it is required to perform SVD of two small-sized matrices R_B and R_W , in $\mathcal{O}(k^3)$ and $\mathcal{O}(n^3)$ flops [16], respectively; as well as to compute a small-sized diagonal matrix exponential in $\mathcal{O}(k)$ flops [17], which are negligible as $k \leq n \ll d$. In the “Computing” phase 6A–10A and 6L–12L, the computational cost is $\mathcal{O}((k+n)d)$. Thus, as $k \leq n \ll d$, the main overhead for solving (3.2) and (3.16) is $\mathcal{O}(d)$.

On the other hand, if the n training vectors are linear independent, the ranks of the matrices S_B and S_W are $k-1$ and $n-k$, respectively, and the main storage requirement of Algorithm 2 is to store $n-1$ vectors of length d . Therefore, the computational complexities of our new algorithms are much fewer than those of Algorithm 1.

3.2. A comparison of the discriminant analysis criteria. In this subsection, we will establish new lower and upper bounds for the criteria (1.4) and (1.6), and give a theoretical comparison of the LDA method and the EDA method. Recall from (3.1) and (3.12) that, the eigenvalues of the matrix M from (3.15) are the same as those of $\exp(-S_W)\exp(S_B)$ and the matrix exponential pencil $(\exp(S_B), \exp(S_W))$.

THEOREM 3.2. Denote by $\nu_1 \geq \nu_2 \geq \dots \geq \nu_d$ the eigenvalues of S_W , by $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$ the eigenvalues of S_B , and let $\lambda_1(M) \geq \lambda_2(M) \geq \dots \geq \lambda_d(M)$ be the eigenvalues of M . Then we have that for $i = 1, 2, \dots, d$,

$$\lambda_i(M) \geq \max \{ \exp(\mu_i - \nu_1), \exp(\mu_d - \nu_{d-i+1}) \}, \quad (3.17)$$

and

$$\lambda_i(M) \leq \min \{ \exp(\mu_i - \nu_d), \exp(\mu_1 - \nu_{d-i+1}) \}. \quad (3.18)$$

Proof. Note that $\lambda_i(M) = \lambda_i(\exp(-S_W)\exp(S_B))$, $i = 1, 2, \dots, d$. Denote $\widetilde{W} = \exp^{1/2}(-S_W)$, we have from the Courant-Fischer minimax theorem that [16]

$$\begin{aligned} \lambda_i(\exp(-S_W)\exp(S_B)) &= \lambda_i(\widetilde{W}^T \exp(S_B) \widetilde{W}) \\ &= \max_{\dim(S)=i} \min_{\mathbf{v} \in S, \|\mathbf{v}\|_2=1} \mathbf{v}^T \widetilde{W}^T \exp(S_B) \widetilde{W} \mathbf{v} \\ &= \max_{\dim(S)=i} \min_{\mathbf{v} \in S, \|\mathbf{v}\|_2=1} \frac{(\widetilde{W} \mathbf{v})^T \exp(S_B) (\widetilde{W} \mathbf{v})}{(\widetilde{W} \mathbf{v})^T (\widetilde{W} \mathbf{v})} \cdot \frac{\mathbf{v}^T \widetilde{W}^T \widetilde{W} \mathbf{v}}{\mathbf{v}^T \mathbf{v}}. \end{aligned} \quad (3.19)$$

Recall that the eigenvalues of $\exp(-S_W)$ and $\exp(S_B)$ are $\exp(-\nu_d) \geq \exp(-\nu_{d-1}) \geq \dots \geq \exp(-\nu_1)$ and $\exp(\mu_1) \geq \exp(\mu_2) \geq \dots \geq \exp(\mu_d)$, respectively. For any $\mathbf{v} \in \mathbb{R}^d$, $\|\mathbf{v}\|_2 = 1$, it follows that

$$\exp(-\nu_1) \leq \frac{\mathbf{v}^T \widetilde{W}^T \widetilde{W} \mathbf{v}}{\mathbf{v}^T \mathbf{v}} = \frac{\mathbf{v}^T \exp(-S_W) \mathbf{v}}{\mathbf{v}^T \mathbf{v}} \leq \exp(-\nu_d), \quad (3.20)$$

and

$$\exp(\mu_i) = \max_{\dim(S)=i} \min_{\mathbf{v} \in S, \|\mathbf{v}\|_2=1} \frac{(\widetilde{W} \mathbf{v})^T \exp(S_B) (\widetilde{W} \mathbf{v})}{(\widetilde{W} \mathbf{v})^T (\widetilde{W} \mathbf{v})}. \quad (3.21)$$

Therefore, we have from (3.19)–(3.21) that

$$\exp(\mu_i - \nu_1) \leq \lambda_i(M) \leq \exp(\mu_i - \nu_d), \quad i = 1, 2, \dots, d. \quad (3.22)$$

On the other hand, we have

$$\lambda_i(\exp(-S_W)\exp(S_B)) = \lambda_i(\exp^{1/2}(S_B)\exp(-S_W)\exp^{1/2}(S_B)).$$

Using the same trick, we can prove that for $i = 1, 2, \dots, d$,

$$\exp(\mu_d - \nu_{d-i+1}) \leq \lambda_i(M) \leq \exp(\mu_1 - \nu_{d-i+1}). \quad (3.23)$$

A combination of (3.22) and (3.23) yields (3.17) and (3.18). \square

We are in a position to establish new lower and upper bounds for the ρ value defined in (1.6). First we need the following lemma.

LEMMA 3.3. [39, Corollary 4.6.4, pp.101] *Let A and B be two $d \times d$ Hermitian matrix, where A is semi-positive definite and B is positive definite. Let X be a $d \times t$ matrix with $\text{rank}(X) = t$. Then*

$$\text{tr}(X^H A X (X^H B X)^{-1}) \geq \sum_{i=1}^t \lambda_{d-t+i}(B^{-1} A),$$

and

$$\text{tr}(X^H A X (X^H B X)^{-1}) \leq \sum_{i=1}^t \lambda_i(B^{-1} A).$$

Combining Theorem 3.2 and Lemma 3.3, we get the theorem as follows for the exponential discriminant analysis criterion (1.6).

THEOREM 3.4. *Under the above notations, for the EDA criterion, there holds*

$$\rho \geq \sum_{i=1}^t \max \{ \exp(\mu_{d-t+i} - \nu_1), \exp(\mu_d - \nu_{t-i+1}) \}, \quad (3.24)$$

and

$$\rho \leq \sum_{i=1}^t \min \{ \exp(\mu_i - \nu_d), \exp(\mu_1 - \nu_{d-i+1}) \}. \quad (3.25)$$

Proof. This result is from Lemma 3.3, (3.17), (3.18), and the fact that for $i = 1, 2, \dots, d$,

$$\lambda_{d-t+i}(\exp(-S_W)\exp(S_B)) \geq \max \{ \exp(\mu_{d-t+i} - \nu_1), \exp(\mu_d - \nu_{t-i+1}) \}.$$

\square

If S_W is nonsingular, we have the following result on the LDA criterion (1.4), whose proof is similar to that of Theorem 3.4, and thus is omitted.

THEOREM 3.5. *If S_W is nonsingular, let $\nu_1 \geq \nu_2 \geq \dots \geq \nu_d > 0$, $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$ be the eigenvalues of S_W and S_B , respectively, and let ϱ be the LDA criterion defined in (1.4). Then we have that*

$$\sum_{i=1}^t \max \left\{ \frac{\mu_{d-t+i}}{\nu_1}, \frac{\mu_d}{\nu_{t-i+1}} \right\} \leq \varrho \leq \sum_{i=1}^t \min \left\{ \frac{\mu_i}{\nu_d}, \frac{\mu_1}{\nu_{d-i+1}} \right\}. \quad (3.26)$$

Without loss of generality, we assume that the n training vectors are linear independent, so that the ranks of the matrices S_B and S_W are $k-1$ and $n-k$, respectively. As $k \leq n \ll d$, there are many 1's in the spectra of $\exp(-S_W)$ and $\exp(S_B)$. The following theorem gives the number of 1 eigenvalues of M .

THEOREM 3.6. *If the n samples in \mathcal{X} are linear independent, then M has at least $d - n + 1$ eigenvalues that are equal to 1.*

Proof. Since the n samples are linear independent, there are $d - k + 1$ and $d - n + k$ zero eigenvalues in the spectra of S_B and S_W , respectively. As a result, there are $d - k + 1$ and $d - n + k$ eigenvalues of $\exp(S_B)$ and $\exp(S_W)$, respectively, which are equal to 1. Let $\{\mathbf{u}_i\}_{i=1}^{d-k+1}$ and $\{\mathbf{g}_i\}_{i=1}^{d-n+k}$ be the orthonormal eigenvectors of $\exp(S_W)$ and $\exp(S_B)$ corresponding to the eigenvalue 1, respectively, and denote by

$$\Phi_1 = \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_{d-k+1}\}, \quad \Phi_2 = \text{span}\{\mathbf{g}_1, \dots, \mathbf{g}_{d-n+k}\}$$

the corresponding eigenspace (or invariant subspace) of $\exp(S_B)$ and $\exp(S_W)$, respectively. Let $\dim(\Phi_i)$ be the dimension of the space Φ_i ($i = 1, 2$), from the well-know dimension formula [39, pp.2]

$$d \geq \dim(\Phi_1 + \Phi_2) = \dim(\Phi_1) + \dim(\Phi_2) - \dim(\Phi_1 \cap \Phi_2),$$

we obtain

$$\dim(\Phi_1 \cap \Phi_2) \geq d - n + 1.$$

Therefore, there are at least $d - n + 1$ independent vectors $\{\mathbf{z}\}_{i=1}^{d-n+1} \in \Omega_1 \cap \Omega_2$, such that

$$\exp(S_W)\mathbf{z}_i = \mathbf{z}_i \quad \text{and} \quad \exp(S_B)\mathbf{z}_i = \mathbf{z}_i.$$

That is,

$$\exp(-S_W)\exp(S_B)\mathbf{z}_i = \mathbf{z}_i, \quad i = 1, 2, \dots, d - n + 1,$$

and M has at least $d - n + 1$ eigenvalues that are equal to 1. \square

REMARK 3.3. *It was assumed in [51, pp.189] that if*

$$\frac{\exp(\mu_i)}{\exp(\nu_i)} > \frac{\mu_i}{\nu_i}, \quad i = 1, 2, \dots, t, \quad (3.27)$$

then there is a difference in diffusion scale between the within-class and between-class distances, and the diffusion scale to the between-class distance is larger than the within-class distance. Hence, the margin between different classes is enlarged, which is helpful in improving classification accuracy. Theorems 3.4–3.6 show this more precisely. It indicates that if

$$\frac{\exp(\mu_i)}{\exp(\nu_d)} > \frac{\mu_i}{\nu_d} \quad \text{or} \quad \frac{\exp(\mu_1)}{\exp(\nu_{d-i+1})} > \frac{\mu_1}{\nu_{d-i+1}}, \quad i = 1, 2, \dots, t, \quad (3.28)$$

then ρ can be larger than ϱ . Notice that the conditions in (3.28) are (much) weaker than those in (3.27). So our result is stronger than the one given in [51], and can be useful to clarify the motivation of EDA.

3.3. Solving the matrix exponential eigenproblems approximately. In the new strategy, one needs to solve the large matrix exponential eigenproblems (3.1) or (3.12) in a prescribed accuracy. If the desired accuracy is high (say, $tol = 10^{-8}$ or even the machine precision $\epsilon \approx 2.22 \times 10^{-16}$), then the cost for solving the large eigenproblems will be very high. In this subsection, we give theoretical analysis on the relationship between the accuracy of the eigenvectors and the distance measure of the nearest neighbour classifier (NN) [10]. Based on the theoretical results, we provide a practical stopping criterion for the matrix exponential eigenproblems.

K nearest neighbors (KNN) is a simple algorithm that stores all available cases and classifies new cases based on a similarity measure [13], e.g., the Euclidean distance. This algorithm has been used in statistical estimation and pattern recognition in the beginning of 1970's as a non-parametric technique. In KNN classification, the output is a class membership. An object is classified by a majority vote of its neighbors, with the object being assigned to the class most common among its K nearest neighbors (where K is a positive integer, typically small). If $K = 1$, then the case is simply assigned to the class of its nearest neighbor (NN) [10].

Denote by $V, \tilde{V} \in \mathbb{R}^{d \times \ell}$ the orthonormal matrices whose columns are the “exact” and “computed” solutions of (3.1) or (3.12), respectively. Let $\hat{\mathbf{x}}_i$ be a sample from the training set, and $\hat{\mathbf{y}}_j$ be a sample from the testing set. Then the nearest neighbour classifier gives class membership via investigating the Euclidean distance [10]

$$d_{ij} = \|V V^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2 = \|V^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2, \quad (3.29)$$

where $\|\cdot\|_2$ denotes 2-norm (or Euclidean norm) of a vector or matrix. The following theorem sheds light on why we can solve the matrix exponential eigenproblems approximately.

THEOREM 3.7. *Let $V, \tilde{V} \in \mathbb{R}^{d \times t}$ be orthonormal matrices whose columns are the “exact” and “computed” solutions of (3.1) or (3.12), respectively. Denote by $d_{ij} = \|V^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2$ and $\tilde{d}_{ij} = \|\tilde{V}^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2$ the “exact” and “computed” Euclidean distances, respectively, and let $\sin \angle(V, \tilde{V}) = \|(I - V V^T)\tilde{V}\|_2$ be the distance between the eigenspace $\text{span}\{V\}$ and the approximate eigenspace $\text{span}\{\tilde{V}\}$. If $\|\hat{\mathbf{x}}_i\|_2, \|\hat{\mathbf{y}}_j\|_2 = 1$ and $\cos \angle(V, \tilde{V}) \neq 0$, then*

$$\frac{\tilde{d}_{ij} - 2 \sin \angle(V, \tilde{V})}{\cos \angle(V, \tilde{V})} \leq d_{ij} \leq \tilde{d}_{ij} \cos \angle(V, \tilde{V}) + 2 \sin \angle(V, \tilde{V}). \quad (3.30)$$

Proof. Note that $V = \tilde{V} \tilde{V}^T V + (I - \tilde{V} \tilde{V}^T) V$, and $\cos \angle(V, \tilde{V}) = \|V^T \tilde{V}\|_2$ [37]. Thus,

$$\begin{aligned} d_{ij} &= \|V^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2 \\ &= \|(V^T \tilde{V}) \tilde{V}^T(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j) + V^T(I - \tilde{V} \tilde{V}^T)(\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j)\|_2 \\ &\leq \tilde{d}_{ij} \|V^T \tilde{V}\|_2 + \|V^T(I - \tilde{V} \tilde{V}^T)\|_2 \cdot \|\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j\|_2 \\ &= \tilde{d}_{ij} \cdot \cos \angle(V, \tilde{V}) + \sin \angle(V, \tilde{V}) \cdot \|\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j\|_2 \\ &\leq \tilde{d}_{ij} \cos \angle(V, \tilde{V}) + 2 \sin \angle(V, \tilde{V}). \end{aligned}$$

where we used $\|\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j\|_2 \leq \|\hat{\mathbf{x}}_i\|_2 + \|\hat{\mathbf{y}}_j\|_2 \leq 2$. Using the same trick, we can prove that

$$\tilde{d}_{ij} \leq d_{ij} \cos \angle(V, \tilde{V}) + 2 \sin \angle(V, \tilde{V}),$$

a combination of the above two inequalities gives (3.30). \square

REMARK 3.4. *Theorem 3.7 indicates that if $\sin \angle(V, \tilde{V})$ is small (say, 10^{-4}), then the $\{d_{ij}\}$'s and the $\{\tilde{d}_{ij}\}$'s will be close to each other. Thus, there is no need to compute the eigenvectors too accurately in practice. This explains why the matrix exponential eigenproblems (3.1) and (3.12) can be solved approximately in practice.*

However, we cannot get the value of $\sin \angle(V, \tilde{V})$ a priori, as the “exact” eigenspace $\text{span}\{V\}$ is unavailable. Let $(\lambda_1, \mathbf{x}_1), \dots, (\lambda_t, \mathbf{x}_t)$ be the desired eigenpairs, it was shown that if the minimal distance between the Ritz values $\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_t$ and the other eigenvalues (i.e., the eigenvalues other than $\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_t$) is sufficiently large, then $\sin \angle(V, \tilde{V})$ is proportional to the residual norms of Ritz pairs [20, 37]. Therefore, we can use the largest residual norm of the Ritz pairs to take the place of $\sin \angle(V, \tilde{V})$ in practice. Experimentally, we find that a tolerance of $\text{tol} = 10^{-4}$ is good enough for the matrix exponential eigenproblems.

In summary, we propose the main algorithm of this work for solving the large generalized matrix exponential eigenproblems.

ALGORITHM 3. Inexact Krylov-EDA algorithms for matrix exponential discriminant analysis

Steps 1–4 are the same as those of Algorithm 2.

5. Given a convergence threshold tol (e.g., $\text{tol} = 10^{-4}$), compute the desired eigenvectors by using a restarted Krylov subspace algorithm (e.g. [3, 36]) for solving (3.1) (by using the Arnoldi method) or (3.13) (by using the Lanczos method), where we use steps 6A–10A or 6L–12L in Algorithm 2 for matrix exponential-vector products.

6. Orthogonalize the columns of the projection matrix V .

REMARK 3.5. *Our new methods fall within the class of the inexact Krylov subspace methods. However, the new method is different from the inner-outer Krylov subspace methods with inexact matrix-vector products [6, 7, 33, 34]. In those methods, the tolerance of the inner iterative process (i.e., approximating the matrix exponential action) can be significantly relaxed once the outer process (i.e., the Arnoldi/Lanczos eigenvalue solver) has started to converge. However, in our new methods the matrix exponential-vector products are computed “exactly” based on some closed-form formulae. Further, the tolerance for Arnoldi/Lanczos eigenvalue solver in Algorithm 3 is determined in advance.*

4. Numerical Experiments. In this section, we make some numerical experiments to illustrate the efficiency of Algorithm 3 for face recognition. All the numerical experiments were run on a Dell PC with eight core Intel(R) Core(TM)i7-2600 processor with CPU 3.40 GHz and RAM 16.0 GB, under the Windows 7 with 64-bit operating system. All the numerical results were obtained from using a MATLAB 7.10.0 implementation. In all the algorithms for comparison, the projection matrix V is composed of the $k - 1$ (i.e., $t = k - 1$) dominant discriminant vectors, where k is the number of classes. We apply the nearest neighbor (NN) [10] as the classifier with the L_2 metric as distance measure. Each column \mathbf{x}_i of the data matrices is scaled by its 2-norm, and 10 random splits are run so that one can obtain a stable recognition result. In all the EDA-based algorithms, the columns of V is orthonormalized by using the MATLAB built-in function `orth.m`, where the `orth` operation is performed via a QR decomposition, i.e., stabilized Gram-Schmidt.

For the performance of the EDA-based methods and how they compare with the current state-of-the-art, we refer to [11, 40, 41, 46, 51]. Indeed, if the large matrix exponential eigenproblems (1.7) and (3.13) were solved “exactly”, then our

Figure 4.1, Example 4.1: Sample face images of 3 individuals from the AR database, $d = 50 \times 40$.

new methods are mathematically equivalent to the original EDA method proposed in [51]. The main aim of this section is to show Algorithm 3 outperforms Algorithm 1 according to CPU time, while the classification accuracy of the former is comparable to that of the latter.

The different EDA-based algorithms are listed as follows:

- **(Inexact) Arnoldi-EDA** (Algorithm 3): We apply Algorithm 3 to solve the large nonsymmetric matrix exponential eigenproblem (3.1), in which the implicitly restarted *Arnoldi* algorithm [36] (the MATLAB built-in function `eigs.m`) is used. The matrix exponential-vector products are computed by using Algorithm 2. The stopping criterion for the large eigenproblem is chosen as $tol = 10^{-4}$.
- **(Inexact) Lanczos-EDA** (Algorithm 3): We apply Algorithm 3 for solving the large symmetric matrix exponential eigenproblem (3.13), in which the implicitly restarted *Lanczos* algorithm [36] (the MATLAB built-in function `eigs.m`) is used. The matrix exponential-vector products are computed by using Algorithm 2. The stopping criterion for the large eigenproblem is chosen as $tol = 10^{-4}$.
- **EDA-eigs**: Solving the generalized symmetric matrix exponential eigenproblem (1.7) by using the MATLAB built-in function `eigs.m`, the convergence threshold is chosen as $tol = 10^{-8}$. This mimics forming the two matrix exponentials $\exp(S_B)$ and $\exp(S_W)$ explicitly while solving the matrix exponential eigenproblem iteratively with a relatively higher accuracy. The matrix exponentials are evaluated by using the MATLAB built-in function `expm.m`.
- **EDA** (Algorithm 1): The EDA algorithm advocated in [51], in which we form $\exp(-S_W)\exp(S_B)$ and solve the nonsymmetric matrix exponential eigenproblem (3.1) by using the implicit QR algorithm [4, 16] (the MATLAB built-in function `eig.m`). The matrix exponentials are evaluated by using the MATLAB built-in function `expm.m`.

Example 4.1 In this example, we compare the EDA-based algorithms with LDA+PCA [5] and the classical LDA method [12, 14], and show the efficiency of the exponential discriminant analysis method. A subset of AR database [21] is used here with 1680 face images from 120 persons (14 images per person), and all images are cropped and scaled to 50×40 . Figure 4.1 presents the sample of cropped AR database images of three individuals. A random subset with $\ell = 2, 3, 5$ images per subject is taken to form the training set, and the rest of the images are used as the testing set. We run Arnoldi-EDA, Lanczos-EDA, EDA-eigs, EDA and LDA+PCA [5] on this problem. As a by-product, we also list the numerical results of the “Classical LDA” method [12, 14], i.e., the classical LDA method with the QZ algorithm [16] for the generalized eigenproblem (1.5) (by using the MATLAB built-in function `eig.m`). As was done in [51], for LDA+PCA, we reserve 99% energy in the PCA stage, followed by LDA. Table 4.1 lists the CPU time (in seconds) and the recognition accuracy obtained from the different methods. When $\ell = 3$, we plot in Figure 4.2 recognition accuracy of the six algorithms with respect to the projected dimension on the AR database. Notice that the curves of Arnoldi-EDA, EDA-eigs, and EDA overlap with each other.

We see from Table 4.1 that the recognition accuracy obtained from the four EDA-based algorithms are about the same, which are (much) higher than those from LDA+PCA and classical LDA. Furthermore, both Arnoldi-EDA and Lanczos-EDA converge much faster than EDA-eigs and EDA, while EDA-eigs outperforms EDA. The classical LDA method works the poorest in terms of CPU time and recognition




Figure 4.2, Example 4.1: Recognition accuracy with respect to the projected dimension on the AR database, 3 Train.

Algorithm	2 Train	3 Train	5 Train
Arnoldi-EDA	0.51(84.3%)	0.71(92.0%)	1.28 (97.0%)
Lanczos-EDA	0.55(84.6%)	0.80(92.2%)	1.58 (97.3%)
EDA-eigs	6.15(84.3%)	6.69(92.0%)	7.35 (97.0%)
EDA	9.61(84.3%)	9.75(92.0%)	10.0 (97.0%)
LDA+PCA	0.09(82.6%)	0.21(87.7%)	0.46 (94.8%)
Classical LDA	51.4(65.0%)	52.1(71.3%)	52.0 (78.0%)

Table 4.1, Example 4.1: CPU time in seconds and recognition accuracy (in brackets) of the six algorithms on the AR database, $d = 50 \times 40$, $k = 120$.

accuracy. This is due to the fact that the dimension $d = 2000 > n = 1680$, and we suffer from the SSS problem. Recall that the classical LDA method cannot cure this drawback properly.

Example 4.2 The aim of this example is two-fold. First, we show that our new algorithms Arnoldi-EDA and Lanczos-EDA run much faster than EDA and EDA-eigs for face recognition. Second, we illustrate the effectiveness of Theorem 3.6. There are two data sets in this example. The first one is the Yale face database taken from the Yale Center for Computational Vision and Control. It contains 165 grayscale images of $k = 15$ individuals. The images demonstrate variation with the following expressions or configurations: 1) lighting: center light, left light, and right light; 2) with/without glasses; and 3) facial expressions: normal, happy, sad, sleepy, surprised, and winking. The original image size is 320×243 pixels. Figure 4.3 gives the sample of cropped Yale database images of three individuals.

The second test set is the Extended YaleB database. This database contains 5760 single light source images of 10 subjects each seen under 576 viewing conditions (9 different poses and 64 illumination conditions of each person). The images have normal, sleepy, sad and surprising expressions. There are some images with or without

Figure 4.3, Example 4.2: Sample face images of 3 individuals from the Yale database, $d = 64 \times 64$.

Figure 4.4, Example 4.2: Sample face images of 3 individuals from the Extended YaleB database, $d = 64 \times 64$.

glasses. These images are captured by varying the position of light source at the center, left or right. For every subject in a particular pose, an image with ambient (background) illumination was also captured. A subset of $k = 38$ with 2432 images are used in this example (64 images of per individual with illumination). Figure 4.4 shows the sample of cropped Extended YaleB database images of three individuals.

In this example, all images are aligned based on eye coordinates and are cropped and scaled to 32×32 , 64×64 and 100×100 , respectively. A random subset with $\ell = 3, 5, 8$ images per subject is taken to form the training set, and the rest of the images are used as the testing set. Tables 4.2 and 4.3 report the CPU time (in seconds) and the recognition accuracy of the different methods.

Three remarks are in order. First, our new algorithms Arnoldi-EDA and Lanczos-EDA outperform EDA-eigs and EDA considerably in terms of CPU time, while the recognition accuracy of the four algorithms are about the same. For example, we see from Table 4.2 that, when $d = 100 \times 100$ and $\ell = 3$, Arnoldi-EDA and Lanczos-EDA used 0.12 and 0.11 seconds, while EDA-eigs and EDA used 409.4 and 1027.4 seconds, respectively. Thus, both the Arnoldi-EDA and the Lanczos-EDA algorithms can circumvent the drawback of heavily computational complexity that bothers the original EDA algorithm, while keep comparable recognition accuracy. Second, cropping the original images may lose some useful information and thus may result in a low recognition accuracy. For instance, for the Yale database, when $\ell = 5$, the recognition accuracy of Arnoldi-EDA applying to $d = 100 \times 100$ images is 93.8%, while the recognition accuracy is only 74.7% as $d = 32 \times 32$. Third, it is seen from Tables 4.2–4.3 that we cannot tell which one, Arnoldi-EDA or Lanczos-EDA, is *definitely better* than the other in terms of CPU time and recognition accuracy.

In order to show effectiveness of Theorem 3.6, we plot in Figure 4.5 the first $k - 1$ eigenvalues and the k -th to the $(d - n + k)$ -th eigenvalues of the matrix $\exp(-S_W)\exp(S_B)$ as $d = 32 \times 32$. One observes that the 4-th to the $(k - 1)$ -th eigenvalues are clustered, and there are at least $d - n + 1$ eigenvalues equal to 1. All these coincide with the theoretical results given in Theorem 3.6.

Example 4.3 In this example, we show efficiency of our “inexact” strategy (see Theorem 3.7 and Remark 3.4) for solving the large matrix exponential eigenproblems (3.1) and (3.13). The FERET database is one of the standard image database specially used for the face recognition algorithms [29]. The final corpus consists of 14051 eight-bit grayscale images of human heads with views ranging from frontal to left and right profiles. In this example, we consider a subset of 1400 images of $k = 200$ individuals, in which each person contributing seven images. The seven images of each individual consists of different illumination and expression variation. Before experiment, the facial images are cropped to a size of 80×80 pixels. Figure 4.6 lists the sample of cropped FERET database images of three individuals. A random subset with $\ell = 2, 3, 5$ images per subject is taken to form the training set, and the rest of the images are used as the testing set. In this example, we have to compute $k - 1 = 199$ dominant eigenpairs of a 6400×6400 matrix using Arnoldi-EDA and Lanczos-EDA, which is a challenging task.




Figure 4.5, Example 4.2: The first $k-1$ eigenvalues (Left) and the k -th to the $(d-n+k)$ -th eigenvalues (Right) of $\exp(-S_W)\exp(S_B)$, Yale data base, $d = 32 \times 32$, $k = 15$; 3 Train.

Algorithm	d	3 Train	5 Train	8 Train
Arnoldi-EDA	32×32	0.04(65.3%)	0.04(74.7%)	0.06(83.1%)
Lanczos-EDA	32×32	0.02(65.2%)	0.02(73.8%)	0.04(83.1%)
EDA-eigs	32×32	0.76(65.8%)	0.79(74.7%)	0.78(83.1%)
EDA	32×32	1.24(65.8%)	1.23(74.7%)	1.24(83.1%)
Arnoldi-EDA	64×64	0.06(74.0%)	0.08(83.8%)	0.11(90.2%)
Lanczos-EDA	64×64	0.04(72.8%)	0.07(84.9%)	0.11(90.2%)
EDA-eigs	64×64	29.7(74.8%)	31.1(84.7%)	31.4(89.8%)
EDA	64×64	73.0(74.8%)	74.7(84.7%)	75.1(89.8%)
Arnoldi-EDA	100×100	0.12(88.3%)	0.15(93.8%)	0.25(96.4%)
Lanczos-EDA	100×100	0.11(88.5%)	0.16(95.3%)	0.31(96.0%)
EDA-eigs	100×100	409.4(88.0%)	397.1(94.8%)	401.4(96.9%)
EDA	100×100	1027.4(88.0%)	993.5(94.9%)	1009.2(96.9%)

Table 4.2, Example 4.2: CPU time in seconds and recognition accuracy (in brackets) of the four algorithms on the Yale database, $k = 15$.

In order to show efficiency of our inexact strategy, we run the (implicitly restarted) Arnoldi and Lanczos algorithm in Step 5 of Algorithm 3 with convergence tolerance $tol = 10^{-2}, 10^{-4}, 10^{-6}, 10^{-8}$ and 10^{-10} , respectively. Table 4.4 presents the numerical results. We observe that the two inexact algorithms work quite well in all the cases, and a tolerance of $tol = 10^{-4}$ is good enough. Indeed, as is shown by Theorem 3.7, it is unnecessary to solve the large exponential eigenproblems in a very high accuracy. This is favorable to very large matrix computations arising from high dimensionality reduction.

It is seen that the recognition rates of Lanczos-EDA are lower than those of Arnoldi-EDA for this problem. A possible reason is that the approximation \mathbf{y} from Lanczos-EDA undergoes an exponential transformation to get the desired solution \mathbf{x} ,

Algorithm	d	3 Train	5 Train	8 Train
Arnoldi-EDA	32×32	0.07(55.9%)	0.11(73.0%)	0.16(82.8%)
Lanczos-EDA	32×32	0.06(59.5%)	0.11(71.2%)	0.15(82.9%)
EDA-eigs	32×32	1.02(56.6%)	1.16(73.2%)	1.28(83.7%)
EDA	32×32	1.28(56.6%)	1.34(73.2%)	1.35(83.7%)
Arnoldi-EDA	64×64	0.23(56.3%)	0.50(72.9%)	0.78(83.2%)
Lanczos-EDA	64×64	0.23(59.4%)	0.58(71.1%)	0.99(83.0%)
EDA-eigs	64×64	36.1(56.9%)	37.7(73.4%)	41.2(84.1%)
EDA	64×64	78.0(56.9%)	78.4(73.4%)	81.4(84.1%)
Arnoldi-EDA	100×100	0.60(56.6%)	1.46(73.9%)	1.90(83.2%)
Lanczos-EDA	100×100	0.74(59.4%)	1.88(71.4%)	2.46(83.3%)
EDA-eigs	100×100	421.7(57.3%)	444.8(73.7%)	541.8(84.2%)
EDA	100×100	1013.2(57.3%)	1030.0(73.7%)	1129.7(84.2%)

Table 4.3, Example 4.2: CPU time in seconds and recognition accuracy (in brackets) of the four algorithms on the Extended YaleB database, $k = 38$.

Figure 4.6, Example 4.3 : Sample face images of 3 individuals from the FERET database, $d = 80 \times 80$.

see (3.14). As a comparison, we also list the numerical results obtained from running EDA and LDA+PCA. One observes that our new algorithms run much faster than EDA, and are superior to LDA+PCA in terms of recognition accuracy.

Algorithm	tol	3 Train	5 Train	8 Train
Arnoldi-EDA	10^{-2}	5.22(62.3%)	6.69(73.4%)	10.2 (81.2%)
	10^{-4}	5.22(62.3%)	6.68(73.3%)	10.3 (81.2%)
	10^{-6}	5.22(62.3%)	10.4(73.4%)	21.6 (81.2%)
	10^{-8}	5.27(62.4%)	11.4(73.4%)	24.6 (81.3%)
	10^{-10}	5.27(62.3%)	12.0(73.4%)	26.8 (81.3%)
Lanczos-EDA	10^{-2}	6.41(44.2%)	8.64(61.7%)	13.8 (79.1%)
	10^{-4}	6.41(44.2%)	8.64(61.7%)	13.8 (79.1%)
	10^{-6}	6.42(44.2%)	14.5(61.6%)	28.0 (79.2%)
	10^{-8}	6.49(44.2%)	15.6(61.6%)	34.9 (79.2%)
	10^{-10}	6.51(44.2%)	16.5(61.6%)	38.0 (79.2%)
EDA	–	288.6(62.3%)	277.7(73.5%)	286.5 (81.4%)
LDA+PCA	–	0.56(42.3%)	1.17(39.7%)	2.54 (57.5%)

Table 4.4, Example 4.3 : CPU time in seconds and recognition accuracy (in brackets) of the Arnoldi-EDA and Lanczos-EDA algorithms (with different tolerances for exponential eigenproblems) on the FERET database, $d = 80 \times 80$, $k = 200$.

Example 4.4 In this example, we demonstrate that Arnoldi-EDA and Lanczos-EDA are suitable to data sets with high dimension. There are two test sets in this example. The ORL database contains 400 face images of $k = 40$ individuals, and the image size is 92×112 . The major challenge on this data set is the variation of the face pose. There is no lighting variation with minimal facial expression variations and no occlusion. Figure 4.7 gives the sample of ORL database images of three individuals. A random subset with $\ell = 2, 3, 5$ images per subject was taken to form the training set, and the rest of the images were used as the testing set.

Figure 4.7, Example 4.4: Sample face images of 3 individuals from the ORL database, $d = 92 \times 112$.

Figure 4.8, Example 4.4: Sample face images of 3 individuals from the CMU-PIE database, $d = 640 \times 486$.

The CMU-PIE database contains more than 40,000 images of 68 subjects with more than 500 images for each. These face images are captured by 13 synchronized cameras and 21 flashes under varying pose, illumination, expression and lights. In this experiment, we choose $k = 10$ subjects and 10 images under different illuminations, lights, expressions and poses for each subject. Thus, the total number of images chosen from CMU-PIE database is 100. The image size is $d = 486 \times 640$ pixels, so we have to deal with eigenproblems of size $311,040 \times 311,040$. We randomly select with $\ell = 3, 5, 8$ images for each subject to organize a training set, while the testing set consists of the remaining images. Figure 4.8 depicts the sample of CUM-PIE database images of three individuals.

In this example, we compare the inexact Krylov subspace algorithms Arnoldi-EDA and Lanczos-EDA with some state-of-the-art algorithms including EDA[51], PCA [38], LDA+PCA (Fisherfaces) [5], NLDA (the null space LDA method) [8], LDA/QR [49], LDA/GSVD [18, 48], RLDA [15] (the regularized LDA method, where the regularized parameter is chosen as 0.01) for the face recognition problem. We reserve 99% energy for recognition for PCA. For LDA+PCA, we also reserve 99% energy in the PCA stage, followed by LDA. Tables 4.5 and 4.6 list the numerical results. In Figures 4.9 and 4.10, we plot curves of recognition accuracy with respect to the projected dimension on the ORL and CUM-PIE databases, with $\ell = 5$ and 3, respectively.

It is observed from Table 4.5 and Figure 4.9 that the recognition accuracies of the four EDA-based algorithms are higher than those of many state-of-the-art algorithms such as LDA+PCA, PCA, LDA/QR and LDA/GSVD. This coincides with the numerical results reported in [11, 40, 41, 46, 51]. We notice that the recognition accuracies of RLDA and NLDA are comparable to those of the four EDA-based methods, while Arnoldi-EDA and Lanczos-EDA run much faster than RLDA, NLDA and LDA/GSVD.

Moreover, it is seen from Table 4.6 and Figure 4.10 that for very large data sets such as CMU-PIE, all the algorithms EDA-eigs, EDA, RLDA, NLDA and LDA/GSVD fail to work, due to their heavy storage requirements. As a comparison, Arnoldi-EDA and Lanczos-EDA perform quite well in all the situations, and their recognition accuracies are a little higher than those of PCA, LDA+PCA and LDA/QR. Consequently, our new algorithms are suitable to data sets with high dimension, which are competitive alternatives to EDA for dimensionality reduction.

5. Conclusions and future work. The computation of large scale matrix exponential eigenproblem is the bottleneck in the frame work of EDA for high dimensionality reduction [1, 11, 40, 41, 46, 51]. In this paper, we propose two inexact Krylov subspace algorithms, i.e., the inexact Arnoldi algorithm and the inexact Lanczos algorithm for nonsymmetric and symmetric matrix exponential eigenproblems, respectively. Our main contribution is to investigate computing matrix exponential-vector products efficiently, and to solve the matrix exponential eigenproblems “approximately”. The relationship between the accuracy of the approximate eigenvectors and the distance to the nearest neighbour classifier (NN) is revealed. A theoretical comparison of the LDA criterion and the EDA criterion is also given. Experimental




Figure 4.9, Example 4.4: Recognition accuracy with respect to projected dimension on the ORL database, 5 Train.




Figure 4.10, Example 4.4: Recognition accuracy with respect to projected dimension on the CMU-PIE database, 3 Train.

results on popular databases for face recognition have validated the effectiveness of the proposed inexact Krylov subspace algorithms.

We point out that our new strategies also apply to other exponential based methods such as the exponential Laplacian embedding, exponential LDE, exponential LPP, exponential discriminant regularization, and 2DEDA for high dimensionality reduction other than face recognition [1, 11, 40, 41, 46, 51]. Furthermore, how to combine our new strategies with other eigen-solvers deserves further investigation. For instance, the Jacobi-Davidson method [35] is a successful eigenvalue solver for large

Algorithm	2 Train	3 Train	5 Train
Arnoldi-EDA	0.36(84.3%)	0.51(87.8%)	0.72 (96.4%)
Lanczos-EDA	0.41(83.1%)	0.63(86.9%)	0.98 (95.7%)
EDA-eigs	407.9(84.3%)	418.2(87.8%)	426.0 (96.4%)
EDA	1079.9(84.3%)	1090.9(87.8%)	1098.9 (96.4%)
LDA+PCA	0.05(83.2%)	0.08(86.7%)	0.17 (91.6%)
PCA	0.03(78.4%)	0.06(83.4%)	0.12 (93.5%)
RLDA	1142.3(84.3%)	1165.7(88.3%)	1159.0 (95.8%)
NLDA	12.0(83.8%)	17.4(87.7%)	16.5 (94.8%)
LDA/QR	0.09(80.2%)	0.12(85.0%)	0.17 (94.0%)
LDA/GSVD	10.9(82.3%)	16.3(85.9%)	15.5 (93.6%)

Table 4.5, Example 4.4: CPU time in seconds and recognition accuracy (in brackets) of the 10 algorithms on the ORL database, $d = 92 \times 112$, $k = 40$.

Algorithm	3 Train	5 Train	8 Train
Arnoldi-EDA	1.98(88.9%)	3.44(94.4%)	5.81(92.0%)
Lanczos-EDA	2.46(90.0%)	4.17(92.4%)	6.93(91.0%)
EDA-eigs	—	—	—
EDA	—	—	—
LDA+PCA	0.36(88.6%)	0.74(94.0%)	1.15(92.0%)
PCA	0.31(82.3%)	0.67(88.4%)	1.06(89.0%)
RLDA	—	—	—
NLDA	—	—	—
LDA/QR	0.47(88.9%)	0.64(93.6%)	0.89(92.0%)
LDA/GSVD	—	—	—

Table 4.6, Example 4.4: CPU time in seconds and recognition accuracy (in brackets) of the 10 algorithms on the CMU-PIE database, $d = 640 \times 486$, $k = 10$. Here “—” means the algorithm suffers from “Out of Memory”.

eigenproblems. This method is applicable to the matrix exponential eigenproblems (3.1) and (3.12). Furthermore, the bound (3.30) only relates the perturbation of the projections to the perturbation of the projected points for the nearest neighbor (NN). It is interesting to establish a performance bound for general KNN which should also involve the distances between points of different classes.

REFERENCES

- [1] N. AHMED. *Exponential discriminant regularization using nonnegative constraint and image descriptor*, IEEE 9th International Conference on Emerging Technologies, pp.1–6, 2013.
- [2] S. AN, W. LIU, S. VENKATESH, AND H. YAN. *Unified formulation of linear discriminant analysis methods and optimal parameter selection*, Pattern Recognition, 44: 307–319, 2011.
- [3] J. BAGLAMA, D. CALVETTI, AND L. REICHEL. *A matlab programe for computing a few eigen pairs of a large sparse Hermitian matrix*, ACM Transactions on Mathematical Software, 29: 337–348, 2003.
- [4] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, Philadelphia, 2000.
- [5] P. BELHUMEUR, J. HESPANHA, AND D. KRIEGMAN. *Eigenfaces vs fisherface: recognition using class-specific linear projection*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 19: 711–720, 1997.
- [6] A. BOURAS AND V. FRAYSSÉ. *A relaxation strategy for the Arnoldi method in eigenproblems*,

- Technical Report TR/PA/00/16, CERFACS, Toulouse, France, 2000.
- [7] A. BOURAS AND V. FRAYSSÉ. *Inexact matrix-vector products in Krylov methods for solving linear systems: a relaxation strategy*, SIAM Journal on Matrix Analysis and Applications, 26: 660–678, 2005.
 - [8] L. CHEN, H. LIAO, M. KO, J. LIN, AND G. YU. *A new lda-based face recognition system which can solve the small-sample-size problem*, Pattern Recognition, 33: 1713–1726, 2000.
 - [9] B. CIPRA. *The best of the 20th century: editors name top 10 algorithms*, SIAM News 33(4), 2000.
 - [10] T. COVER AND P. HART. *Nearest neighbor pattern classification*, IEEE Transactions on Information Theory, 13: 21–27, 1967.
 - [11] F. DORNAIKA AND A. BOSAGHZADEH. *Exponential local discriminant embedding and its application to face recognition*, IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics, 43: 921–934, 2013.
 - [12] R. DUDA, P. HART, AND D. STORK. *Pattern Classification*, 2nd ed. New York: Wiley, 2000.
 - [13] B. EVERITT, S. LANDAU, M. LEESE, AND D. STAHL. *Miscellaneous Clustering Methods*, in Cluster Analysis, 5th Edition, John Wiley & Sons, Ltd, Chichester, UK, 2011.
 - [14] R. FISHER. *The statistical utilization of multiple measurements*, Annals and Eugenics, 8: 376–386, 1938.
 - [15] J. FRIEDMAN. *Regularized discriminant analysis*, Journal of the America Statistical Association, 84: 165–175, 1989.
 - [16] G.H. GOLUB AND C.F. VAN LOAN. *Matrix Computations*. The Johns Hopkins university Press, 4th edition, 2013.
 - [17] N.J. HIGHAM. *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008.
 - [18] P. HOWLAND AND H. PARK. *Generalizing discriminant analysis using the generalized singular value decomposition*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 26: 995–1006, 2004.
 - [19] Y. JIA, F. NIE, AND C. ZHANG. *Trace ratio problem revisited*, IEEE Transactions on Neural Networks, 20: 729–735, 2009.
 - [20] Z. JIA AND G.W. STEWART. *An analysis of the Rayleigh-Ritz method for approximating eigenspaces*, Mathematics of Computation, 70: 637–647, 2001.
 - [21] A. MARTINEZ. *Recognizing imprecisely localized, partially occluded, and expression variant faces from a single sample perclass*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 24: 748–763, 2002.
 - [22] C. MOLER AND C.F. VAN LOAN. *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Review, 45: 3–49, 2003.
 - [23] T. NGO, M. BELLALI, AND Y. SAAD. *The trace optimization problem for dimensionality reduction*, SIAM Review, 54: 545–569, 2012.
 - [24] F. NIE, S. XIANG, Y. JIA, C. ZHANG, AND S. YAN. *Trace ratio criterion for feature selection*. Proceedings of the 23rd AAAI Conference on Artificial Intelligence, pp.671–676, 2008.
 - [25] F. NIE, S. XIANG, Y. JIA, AND S. ZHANG. *Semi-supervised orthogonal discriminant analysis via label propagation*, Pattern Recognition, 42: 2615–2627, 2009.
 - [26] F. NIE, S. XIANG, Y. LIU, C. HOU, AND C. ZHANG. *Orthogonal vs. uncorrelated least squares discriminant analysis for feature extraction*, Pattern Recognition Letters, 33: 485–491, 2012.
 - [27] F. NIE, S. XIANG, AND C. ZHANG. *Neighborhood MinMax Projections*, The Twentieth International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India, pp. 993–997, 2007.
 - [28] C. PARK AND H. PARK. *A comparison of generalized linear discriminant analysis algorithms*, Pattern Recognition, 41: 1083–1097, 2008.
 - [29] P. PHILLIPS, H. MOON, P. J. RAUSS, AND S. RIZVI. *The FERET evaluation methodology for face recognition algorithms*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 22: 1090–1104, 2000.
 - [30] M. POPOLIZIO AND V. SIMONCINI. *Acceleration techniques for approximating the matrix exponential operator*, SIAM Journal on Matrix Analysis and Applications, 30: 657–683, 2008.
 - [31] Y. SAAD. *Numerical Methods for Large Eigenvalues Problems*, 2nd edition, SIAM, Philadelphia, 2011.
 - [32] Y. SAAD. *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM Journal on Numerical Analysis, 29: 209–228, 1992.
 - [33] V. SIMONCINI. *Variable accuracy of matrix-vector products in projection methods for eigencomputation*, SIAM Journal on Numerical Analysis, 43: 1155–1174, 2005.
 - [34] V. SIMONCINI AND D. SZYLD. *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM Journal on Scientific Computing, 25: 454–477, 2003.

- [35] G. SLEIJPEN AND H. VAN DER VORST. *A Jacobi-Davidson iteration method for linear eigenvalue problems*, SIAM Journal on Matrix Analysis and Applications, 17: 401–425, 1996.
- [36] D. SORESENSEN. *Implicit application of polynomial filters in a k -step Arnoldi method*, SIAM Journal on Matrix Analysis and Applications, 13: 357–385, 1992.
- [37] G.W. STEWART. *Matrix Algorithms II: Eigensystems*, SIAM, Philadelphia, 2001.
- [38] M. TURK AND A. PENTLAND. *Eigenfaces for recognition*, Cognitive Neuroscience, 3: 71–86, 1991.
- [39] S.G. WANG, M. WU, AND Z.Z. JIA. *Matrix inequalities* (in Chinese), Science Academic Press, Beijing, 2nd edition, 2005.
- [40] S. WANG, H. CHEN, X. PENG, AND C. ZHOU. *Exponential locality preserving projections for small sample size problem*, Neurocomputing, 74: 36–54, 2011.
- [41] S. WANG, S. YAN, J. YANG, C. ZHOU, AND X. FU. *A general exponential framework for dimensionality reduction*, IEEE Transactions on Image Processing, 23: 920–930, 2014.
- [42] X. WANG AND X. TANG. *A unified framework for subspace face recognition*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 26: 1222–1228, 2004.
- [43] G. WU AND Y. WEI. *On analysis of projection methods for rational function approximation to the matrix exponential*, SIAM Journal on Numerical Analysis, 48: 191–197, 2010.
- [44] G. WU, W. XU, AND H. LENG. *Inexact and incremental bilinear Lanczos components algorithms for high dimensionality reduction and image reconstruction*, Pattern Recognition, 48: 244–263, 2015.
- [45] G. WU, H. PANG, AND J. SUN. *Preconditioning the restarted and shifted block FOM algorithm for matrix exponential computation*, arXiv preprint, available online at <http://arxiv.org/abs/1405.0707>.
- [46] L. YAN AND J. PAN. *Two-dimensional exponential discriminant analysis and its application to face recognition*, International Conference on Computational Aspects of Social Networks (CASoN), pp. 528–531, 2010.
- [47] J. YE. *Charaterization of a family of algorithms for generalized discriminant analysis on undersampled problems*, Journal of Machine Learning Research, 6: 483–502, 2006.
- [48] J. YE, R. JANARDAN, C.H. PARK, AND H. PARK. *An optimization criterion for generalized discriminant analysis on undersampled problems*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 26: 982–994, 2004.
- [49] J. YE AND Q. LI. *A two-stage linear discriminant analysis via QR decomposition*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 27: 929–941, 2005.
- [50] H. YU AND H. YANG. *A direct LDA algorithm for high-dimensional data with application to face recognition*, Pattern Recognition, 34: 2067–2070, 2001.
- [51] T. ZHANG, B. FANG, Y. TANG, Z. SHANG, AND B. XU. *Generalized discriminant analysis: a matrix exponential approach*, IEEE Transactions on Systems Man and Cybernetics-part B: cybernetics, 40: 186–197, 2010.